

Flexible Loop Filter Design for Spacecraft Phase-Locked Receivers

J. B. Berner, J. M. Layland, and P. W. Kinman

Abstract

A flexible loop filter design for spacecraft phase-locked receivers is proposed. The loop filter is implemented as digital hardware with coefficients that are set by registers. Either a perfect or an imperfect integrating loop filter can be effected. This flexibility is important since no one type of loop filter is preferred for all circumstances. An imperfect integrator is preferred when, as is often the case for spacecraft receivers, it is important to minimize the best-lock frequency drift of an idling loop. A perfect integrator is preferred when the tracking performance of the loop is the most important consideration.

1 Introduction

Phase-locked loop receivers are used on spacecraft to provide carrier synchronization of the uplink. This is an essential function for the demodulation of phase-shift keyed command channels and for two-way coherent Doppler and range measurement. Such loops are usually of second order, since this provides the best compromise between tracking performance and stability for a space-based receiver. A second-order phase-locked loop may have either a perfect or an imperfect integrating loop filter. The former gives the better tracking performance. The latter is subject to less best-lock frequency drift while idling (that is, operating with no signal, only noise, at the input). For many space missions it is important that the best-lock frequency not drift too much while the carrier loop idles so that acquisition is quick when a new uplink arrives. In short, there is a trade-off in the selection of loop filter type. Even within a given space mission, we might want to change the type of loop filter. For those phases of the mission for which the Doppler dynamics are challenging we prefer a perfect integrating loop filter, optimizing tracking performance, and for other phases of the mission during which there are frequency interruptions in radio contact we prefer an imperfect integrating loop filter, minimizing best-lock frequency drift and speeding acquisition.

The loop filter is nowadays typically implemented in digital hardware. We propose a loop filter structure that is implemented in digital hardware with coefficients that are set by registers. The design is flexible enough to accommodate either type of loop filter. Most important, the loop may be changed from one type to the other by simply rewriting the coefficients.

We analyze in this paper the best-lock frequency drift of an idling phase-locked loop receiver due to the presence of Gaussian (thermal and shot) noise at its input and show that a receiver with an imperfect integrating loop filter drifts less than one of the same loop bandwidth with a perfect integrating loop filter. (We do not analyze tracking performance of a phase-locked receiver because that is well documented in the literature.) Then we introduce the flexible loop filter design that accommodates both a perfect and an imperfect integrating loop filter. We begin by modeling the signal plus noise in a digital phase-locked loop receiver; this preliminary work is necessary for the analysis of best-lock frequency drift.

2 Modeling of a Digital Phase-Locked Loop

We model a digital phase-locked loop with the signal processing of Figure 1. Other equivalent structures are possible—a loop implemented with synchronous sampling, for example. In this

paper, we think in terms of Figure 1 only to make the discussion specific; but we bear in mind that equivalent loop architectures are subject to the same set of considerations.

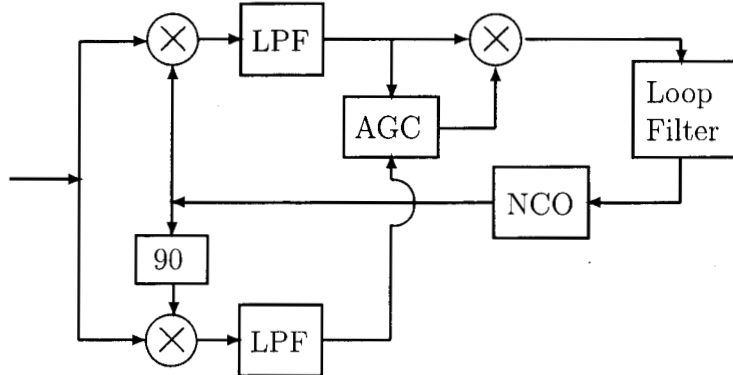


Figure 1: Digital phase-locked loop with AGC.

The input to the digital phase-locked loop is of the form

$$\sqrt{2P} \sin[\omega k + \theta(k)] + n(k).$$

We assume that the sampling is periodic. We denote by T the sample period and by k the discrete-time index. The nominal phase advance (modulo 2π radians) of the carrier between adjacent samples is denoted ω (in radians). Any deviation from the nominal phase advance is accounted for in the phase term $\theta(k)$. The power in the received carrier is denoted P .

We model the input noise $n(k)$ as a sequence of independent noise samples, each Gaussian with a mean of zero and a variance $N_0/(2T)$. The Gaussian distribution is appropriate for thermal and shot noise, which typically constitute the dominant noise of spacecraft receivers. We do not consider phase noise, such as might originate in the transmitter or be picked up in the medium, which might have different statistics. The modeling of the noise as uncorrelated from one sample to the next may seem unrealistic. In principle, a white noise sequence like this results when the sampling is preceded by an anti-aliasing filter that has a transfer function with a rectangular shape and a bandwidth equal to one-half the sampling rate. In practice, of course, no causal analog filter has a rectangular transfer function. A practical anti-aliasing filter has a bandwidth that is somewhat smaller than one-half the sampling rate and has finite roll-off in its transfer function. Thus, in a practical receiver, the noise would have some correlation from one sample to the next and would have a variance somewhat smaller than the $N_0/(2T)$ that we use in our model. It will be noted that our simple noise model is pessimistic in that it uses a too-large variance and is optimistic in that it assumes uncorrelated noise samples. Experience shows that the pessimism approximately cancels the optimism. Within the loop there is low-pass filtering, which comes in the form of averaging, and this additional filtering greatly reduces the contribution of that part of the noise lying in the frequency domain close to the band edges of the anti-aliasing filter. So that part of the noise that we have inaccurately modeled is later filtered out; the fact that our modeling of it was flawed becomes inconsequential. In summary, then, we use what may seem like an over-simplified model for noise, but this model leads to good, approximate results.

The output of the Numerically-Controlled Oscillator (NCO) is modeled as

$$\sqrt{2} \cos[\omega k + \hat{\theta}(k)].$$

The output of the upper (quadrature) multiplier is of the form

$$\sqrt{P} \sin \phi(k) + n_Q(k).$$

The loop phase error is denoted

$$\phi(k) = \theta(k) - \hat{\theta}(k). \quad (1)$$

There is also a sum-frequency signal component, but it is subsequently filtered out, and so we feel free to ignore it here. The quadrature noise $n_Q(k)$ is a sequence of zero-mean, Gaussian, independent noise samples, each of variance $N_0/(2T)$.

The local oscillator input to the lower (in-phase) multiplier comes from the NCO, but with a phase delay of 90° ; this local oscillator is of the form

$$\sqrt{2} \sin [\omega k + \hat{\theta}(k)].$$

The output of the lower (in-phase) multiplier is of the form

$$\sqrt{P} \cos \phi(k) + n_I(k).$$

The in-phase noise $n_I(k)$ is a sequence of zero-mean, Gaussian, independent noise samples, each of variance $N_0/(2T)$. The cross-correlation between in-phase and quadrature noise is very small.

A typical spacecraft receiver must operate over a large dynamic range. In order that the loop bandwidth and loop damping factor not vary wildly, we must control the signal level within the loop. In practice, we remove most, but not all, of the signal-level variation using Automatic Gain Control (AGC). This is accomplished by the signal processing described here. The quadrature samples are averaged, M samples at a time; this has the effect of multiplying the noise variance by the factor $1/M$. The taking of an average is a kind of low-pass filtering; and so, of course, the noise variance decreases. In parallel, the in-phase samples are averaged, M samples at a time. The sum of the squares of these two parallel averages is computed. This sum has an expected value of $P + N_0/(MT)$. This sum is further averaged in order to get a power indicator with a small deviation. The reader may well wonder why we don't do *all* the averaging *before* the nonlinear operation of squaring. If we did that, we would remove *all* the signal-level variation. As will be explained in more detail later, it is advantageous to leave some signal-level variation because it gives the receiver a useful adaptive capability—the receiver can automatically narrow the loop bandwidth in response to an unusually weak signal. The AGC signal processing then computes a factor G as the reciprocal square-root of the averaged sum. That is to say,

$$G = \frac{1}{\sqrt{P + \frac{N_0}{MT}}}. \quad (2)$$

The output of the quadrature multiplier is then scaled by this factor G (using the second multiplier in the quadrature branch of Figure 1). The result of this scaling, which is the input to the loop filter, is denoted $x(k)$ and is of the form

$$x(k) = \alpha \sin \phi(k) + n_G(k). \quad (3)$$

The suppression factor α is given by

$$\alpha = \frac{1}{\sqrt{1 + \frac{N_0}{PMT}}}. \quad (4)$$

The noise $n_G(k)$ at the input to the loop filter is a sequence of zero-mean, Gaussian, independent noise samples, each of variance $G^2 \cdot N_0/(2T) = (1 - \alpha^2)M/2$.

When the loop is in phase-lock, $\phi(k) \ll 1$ radian, and $\sin \phi(k) \approx \phi(k)$. Under these circumstances, the incoming phase $\theta(k)$ and the tracking phase $\hat{\theta}(k)$ have a linear, time-invariant relationship, as depicted in Figure 2. From this figure, it can be seen that the tracking phase $\hat{\theta}(k)$ can be thought of as the low-pass filtering of the sum

$$\theta(k) + \frac{1}{\alpha} \cdot n_G(k),$$

where the low-pass filter has a (discrete-time) transfer function given by

$$\frac{\alpha F(z)}{(z-1) + \alpha F(z)}.$$

The loop filter has a (discrete-time) transfer that we denote $F(z)$, the form of which is discussed below. The scaled noise $n_G(k)/\alpha$ has variance $N_0/(2PT)$.

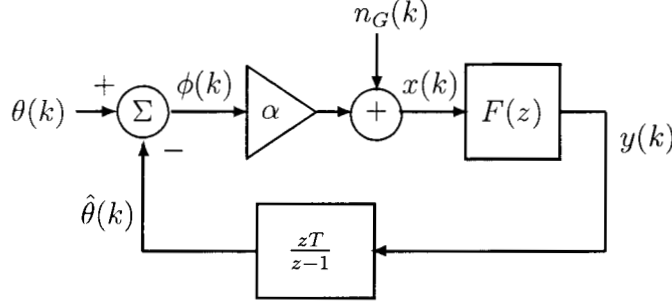


Figure 2: Approximate linear, time-invariant relationship between $\theta(k)$ and $\hat{\theta}(k)$ during phase-lock.

Two types of loop filter are considered in this paper, one with perfect integrator and one with imperfect integrator. The corresponding loop filter (discrete-time) transfer functions are

$$F(z) = \begin{cases} K_1 z^{-1} + K_2 \cdot \frac{T}{z-1} & \text{perfect integrator} \\ K z^{-1} \cdot \frac{T + \tau_2(z-1)}{T + \tau_1(z-1)} & \text{imperfect integrator.} \end{cases} \quad (5)$$

We now want to evaluate the transfer function of the low-pass filter that characterizes the relationship of Figure 2. This is most easily done if we use the continuous-update approximation. [1] This approximation follows from the substitution

$$\frac{z-1}{T} \rightarrow s,$$

so that we remove the z -transform variable and insert the Laplace transform variable s . In addition, we ignore z wherever it occurs outside of the difference $z-1$. This continuous-update approximation is valid whenever the product of T and the loop bandwidth is less than about 0.1. With this approximation the phase-locked loop transfer function becomes

$$H(s) = \begin{cases} \frac{\alpha K_1 s + \alpha K_2}{s^2 + \alpha K_1 s + \alpha K_2} & \text{perfect integrator} \\ \frac{\alpha K \tau_2 s + \alpha K}{\tau_1 s^2 + (1 + \alpha K \tau_2) s + \alpha K} & \text{imperfect integrator.} \end{cases} \quad (6)$$

For both types of loop filters considered in this paper, the loop transfer function is second-order, which is typical of spacecraft receivers. The noise-equivalent loop bandwidth B may be found from $H(s)$. Since $H(0) = 1$,

$$B = \int_0^\infty |H(j2\pi f)|^2 df. \quad (7)$$

Eq. (7) has been evaluated for both perfect integrating loop filters [1] and imperfect integrating loop filters [2]:

$$B = \begin{cases} \frac{\alpha K_1^2 + K_2}{4K_1} & \text{perfect integrator} \\ \frac{\alpha K (\tau_1 + \alpha K \tau_2^2)}{4\tau_1 (\alpha K \tau_2 + 1)} & \text{imperfect integrator.} \end{cases} \quad (8)$$

Another important parameter for a second-order control loop is the damping factor ζ ,

$$\zeta = \begin{cases} \frac{K_1}{2} \sqrt{\frac{\alpha}{K_2}} & \text{perfect integrator} \\ \frac{1 + \alpha K \tau_2}{2\sqrt{\alpha K \tau_1}} & \text{imperfect integrator.} \end{cases} \quad (9)$$

By way of example, four specific phase-locked loops are considered in this paper. They are listed in Table 1. Loops L_1 and L_3 have perfect integrator loop filters. Loops L_2 and L_4 have imperfect integrator loop filters. For all loops considered in this paper, we use $M = 8$ and $T^{-1} = 75,000 \text{ s}^{-1}$.

Both B and ζ depend on the signal-to-noise spectral density ratio P/N_0 . This is because they depend on α and α is a function of P/N_0 , as seen in Eq. (4). Figure 3 plots B and Figure 4 plots ζ as a function of P/N_0 for the loops of Table 1. For loops L_1 and L_2 , the loop bandwidth is approximately the same function of P/N_0 ; for large values of P/N_0 , the asymptote is $B = 90 \text{ Hz}$. For loops L_3 and L_4 , the loop bandwidth is approximately the same function of P/N_0 ; for large values of P/N_0 , the asymptote is $B = 200 \text{ Hz}$. The loop damping factor is approximately the same function of P/N_0 for all four loops.

Loop	Integration	Parameters
L_1	perfect	$K_1 = 342$ $K_2 = 6190$
L_2	imperfect	$K = 2.2 \times 10^7$ $\tau_1 = 3556$ $\tau_2 = 0.0556$
L_3	perfect	$K_1 = 760$ $K_2 = 30600$
L_4	imperfect	$K = 3.0 \times 10^7$ $\tau_1 = 1000$ $\tau_2 = 0.025$

Table 1: The four loops considered in this paper.

There is always a trade-off in selecting a loop bandwidth. With a small B we minimize the effect of noise on our loop, and with a large B we get better tracking of the incoming signal phase. (Recall that the tracking phase is essentially a low-pass filtering of the incoming signal phase plus noise.) For the phase-locked carrier loop of a spacecraft receiver, it is desirable to have the kind of loop bandwidth variation shown in Figure 3. [3] For large P/N_0 , we want a large B in order to get good signal tracking. For a small P/N_0 , we want a relatively small B so that the signal-to-noise ratio in the loop is maximized for the available P/N_0 . The received P/N_0 might vary over several tens of decibels during the life of the spacecraft, but we don't want B to also vary by several orders of magnitude. Instead, we seek a milder variation in B ; and this is achieved with the AGC signal processing already described for a judicious choice for the product MT . The variation in ζ is incidental; it accompanies the variation in B . The exact value of ζ is never important, but typically we like it not to stray too far from 1.

For a given loop bandwidth, a perfect integrator results in better tracking performance than does an imperfect integrator. [4] For example, when the arriving signal has a constant frequency offset δf from the best-lock frequency of the phase-locked loop, a perfect integrator results in zero phase error but an imperfect integrator gives a phase error of $2\pi\delta f/(\alpha K)$, where αK is the loop gain.

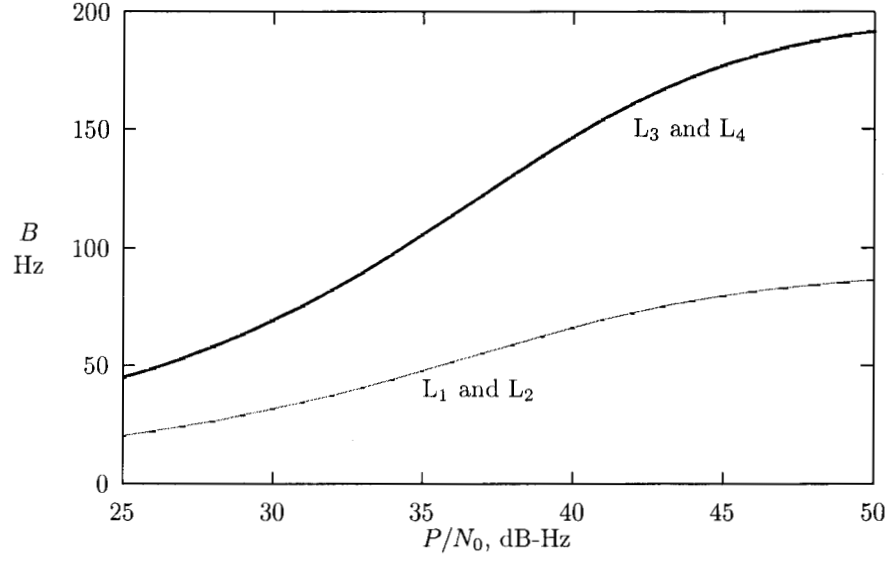


Figure 3: Noise-equivalent loop bandwidth; $M = 8$, $T^{-1} = 75,000 \text{ s}^{-1}$.

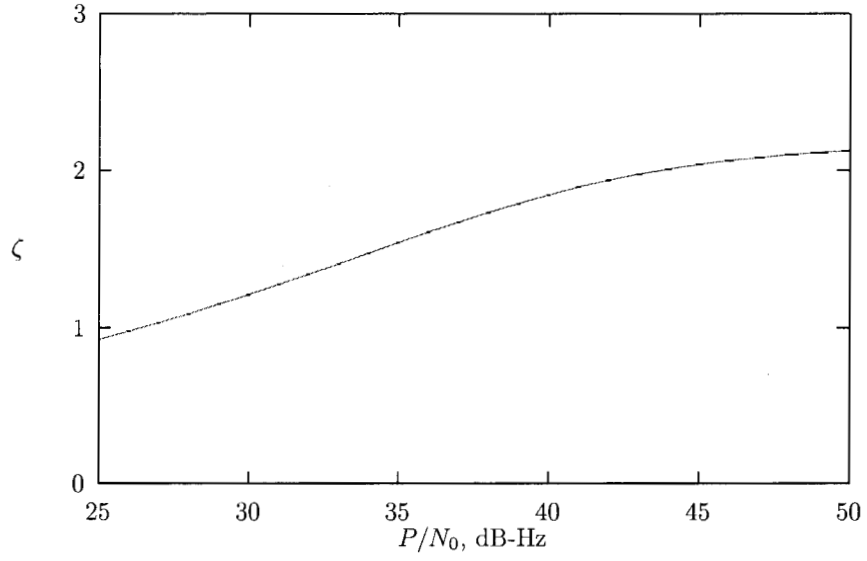


Figure 4: Loop damping factor; $M = 8$, $T^{-1} = 75,000 \text{ s}^{-1}$.

3 Best-Lock Frequency Drift in an Idling Loop

It should not be concluded, though, that it is always better to use a perfect integrator. There is one respect in which a loop with imperfect integrator is better. When the loop idles (i.e., when noise alone is at the input) and for a given loop bandwidth, the best-lock frequency of the loop will drift less in response to noise at its input if the loop filter integration is imperfect, rather than perfect. Spacecraft receivers sometimes must idle for many hours during a solar conjunction or while waiting for a ground tracking station to become available. When a direct line-of-sight radio connection again becomes possible between a ground-based transmitter and the spacecraft receiver, it is often important that the receiver acquire phase-lock as quickly as possible. To get quick acquisition, we must know the best-lock frequency of the spacecraft receiver to good accuracy. Therefore, we would like the best-lock frequency not to have drifted much since the end of the previous communications session.

In this section of the paper, we investigate the drift of best-lock frequency caused by thermal noise acting on an idling loop. It should be mentioned that this is not the only cause of best-lock frequency drift; there will always be some drift due to temperature variation of the receiver crystal oscillators.

When the receiver idles, so that noise alone is at the loop input, $P = 0$, and this implies $\alpha = 0$. The loop is effectively open. The input $x(k)$ to the loop filter is, from Eq. (3), $n_G(k)$, whose variance is in this case just $M/2$.

When the loop filter is a perfect integrator, it is possible to calculate the root-mean-square (rms) best-lock frequency offset σ_f due to noise at the input of the loop filter. In this case, the loop filter output $y(k)$ is related to $x(k)$ by

$$y(k) = K_1 x(k-1) + K_2 T \sum_{i=0}^{k-1} x(i). \quad (10)$$

Loss-of-lock corresponds to $k = 0$. Substituting $n_G(k)$ (with a variance of $M/2$) for $x(k)$ gives a variance σ_y^2 on $y(k)$ of

$$\sigma_y^2 = (K_1^2 + 2K_1 K_2 T + k^2 K_2^2 T^2) \frac{M}{2}. \quad (11)$$

Since $y(k)$ is an angular frequency in units of rad/s, σ_f can be found from

$$\sigma_f = \frac{\sigma_y}{2\pi} = \frac{1}{2\pi} \sqrt{(K_1^2 + 2K_1 K_2 T + k^2 K_2^2 T^2) \frac{M}{2}}. \quad (12)$$

Denoting the time since loss-of-lock as t ,

$$t = kT, \quad (13)$$

we can rewrite Eq. (12) as

$$\sigma_f = \frac{\sigma_y}{2\pi} = \frac{1}{2\pi} \sqrt{(K_1^2 + 2K_1 K_2 T + K_2^2 t^2) \frac{M}{2}}. \quad (14)$$

When the loop filter is an imperfect integrator, no such simple analytical solution for σ_f exists, but the action of the loop filter is easily simulated. The loop filter output $y(k)$ and input $x(k)$ are related by the difference equation

$$\tau_1 y(k) + (T - \tau_1) y(k-1) = K \tau_2 x(k-1) + K(T - \tau_2) x(k-2). \quad (15)$$

As before, $x(k) = n_G(k)$ with a variance of $M/2$.

The rms best-lock frequency offset σ_f is plotted in Figure 5 for loops L_1 and L_2 and in Figure 6 for loops L_3 and L_4 . For the loops with perfect integrator (L_1 and L_3), Eq. (14) was used. For the loops with imperfect integrator (L_2 and L_4), simulation results are plotted. From either of these figures, we see that for two loops with comparable bandwidths, the imperfect integrator

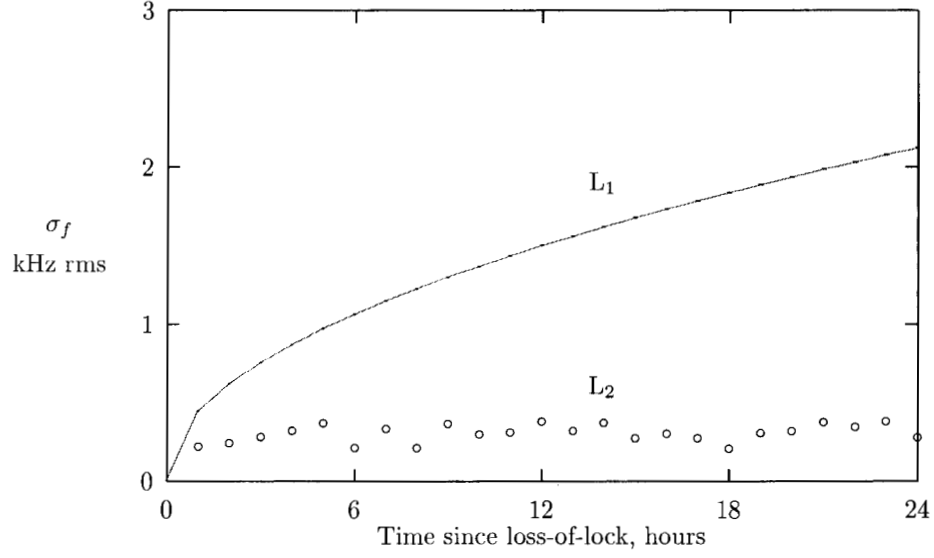


Figure 5: Best-lock frequency drift for loops L_1 and L_2 ; $M = 8$, $T^{-1} = 75,000 \text{ s}^{-1}$.

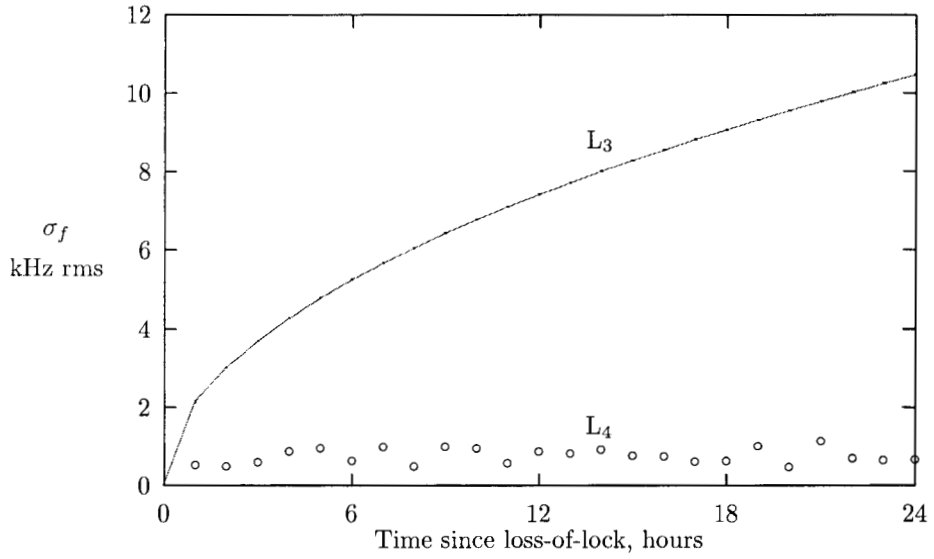


Figure 6: Best-lock frequency drift for loops L_3 and L_4 ; $M = 8$, $T^{-1} = 75,000 \text{ s}^{-1}$.

loop experiences less drift of the best-lock frequency in response to noise while idling than does the perfect integrator loop. In comparing Figures 5 and 6, we see that a loop of a given type experiences less drift of the best-lock frequency in response to noise if it has a smaller loop bandwidth.

In principle, it is possible to reduce the best-lock frequency drift due to noise by decreasing M (the number of samples that are averaged, as part of the AGC signal processing, in each of the quadrature channels before the sum-of-squares operation). But decreasing M can result in a too-strong dependence of B and ζ on P/N_0 . As explained in the previous section, we want B and ζ to have a mild dependence on P/N_0 , as typified by Figure 3. A stronger dependence on P/N_0 would mean, for example, that ζ is far from 1.0 (critical damping) much of the time. (A ζ that is close to 1.0 results in good loop transient response.) So decreasing M is not really an option.

4 Flexible Loop Filter Implementation

The previous sections have demonstrated that the choice between perfect and imperfect integrator depends on the relative importance of minimizing best-lock frequency drift and maximizing tracking performance. For a given space mission, we may want a perfect integrating loop filter for some phases of the mission, during which tracking performance is the paramount criterion, and for other phases we may want an imperfect integrating loop filter to ensure fast acquisition following periods of idling.

We propose a loop filter structure for implementation in digital hardware, featuring coefficients set by registers, that can be configured for either perfect or imperfect integration. Figure 7 shows this structure. In that figure, the final delay element represents the transport delay of the loop. The flexibility arises from the fact that the coefficients are set by registers and, therefore, can be altered even after the launch of the spacecraft.

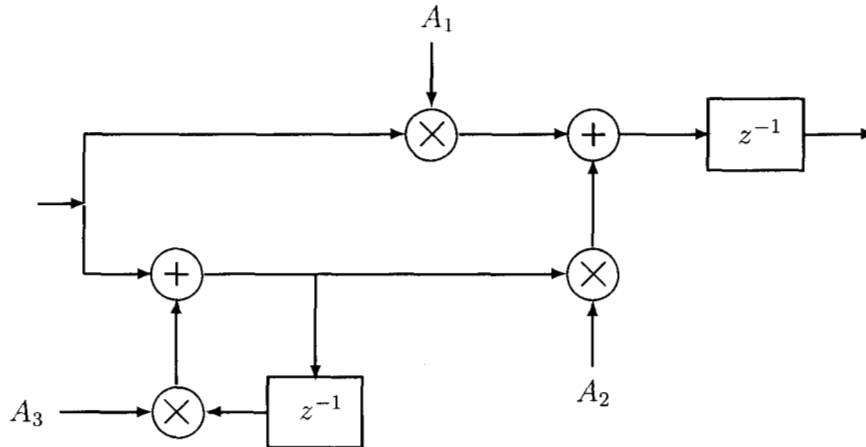


Figure 7: Loop filter implementation.

The transfer function of the loop filter shown in Figure 7 is given by

$$A_1 z^{-1} + \frac{A_2}{z - A_3}.$$

This transfer function represents a perfect integrator when

$$\left. \begin{aligned} A_1 &= K_1 \\ A_2 &= K_2 T \\ A_3 &= 1. \end{aligned} \right\} \quad (16)$$

Alternatively, the general transfer function for the implementation of Figure 7 may be put in the form

$$z^{-1} \left[\frac{\tau_1(A_1 + A_2 - A_1 A_3) + \tau_1(A_1 + A_2)(z - 1)}{\tau_1(1 - A_3) + \tau_1(z - 1)} \right].$$

This matches the form of an imperfect integrator when

$$\left. \begin{aligned} \tau_1(1 - A_3) &= T \\ \tau_1(A_1 + A_2) &= K\tau_2 \\ \tau_1(A_1 + A_2 - A_1 A_3) &= KT. \end{aligned} \right\} \quad (17)$$

Solving for the parameters A_1 , A_2 and A_3 gives

$$\left. \begin{aligned} A_1 &= K \frac{T - \tau_2}{T - \tau_1} \\ A_2 &= K \left[\frac{\tau_2}{\tau_1} - \frac{T - \tau_2}{T - \tau_1} \right] \\ A_3 &= 1 - \frac{T}{\tau_1}. \end{aligned} \right\} \quad (18)$$

Hence, the one loop filter structure shown in Figure 7 produces a perfect integrator when the parameters A_1 , A_2 and A_3 are assigned the values given in Eqs. (16). When, instead, the parameters are assigned values according to Eqs. (18), an imperfect integrator is effected.

Table 2 shows the appropriate values of A_1 , A_2 and A_3 for each of the four loops of Table 1. These values are for the case of $T^{-1} = 75,000 \text{ s}^{-1}$.

A few comments are in order about the precision with which the parameters should be set. A precision of about $\pm 1\%$ in the setting of A_1 and A_2 will result in a precision of about $\pm 1\%$ on B . The loop design is much more sensitive to the value of A_3 . For the loops L_1 and L_3 , which have perfect integrator loop filters, A_3 should equal *exactly* 1. For the loops L_2 and L_4 , which have imperfect integrator loop filters, A_3 needs to equal 1 minus a small, precise amount,

$$A_3 = 1 - \epsilon, \quad (19)$$

where

$$\epsilon = \begin{cases} 3.750 \times 10^{-9} & L_2 \\ 1.333 \times 10^{-8} & L_4. \end{cases} \quad (20)$$

ϵ should be set with a precision of about $\pm 1\%$ in order to get a precision of about $\pm 1\%$ on B .

Loop	A_1	A_2	A_3
L_1	342.0	0.0825	1.00000000000
L_2	343.9	0.0825	0.99999999625
L_3	760.0	0.4080	1.00000000000
L_4	749.6	0.4000	0.99999998666

Table 2: A_1 , A_2 and A_3 for the loops of Table 1 with $T^{-1} = 75,000 \text{ s}^{-1}$.

5 Conclusions

The best-lock frequency of an idling phase-locked loop will drift in response to noise at the receiver input. This drift is worse for a loop with perfect integrator than it is for one with imperfect integrator. On the other hand, a loop with perfect integrator offers better tracking performance. We have proposed a flexible loop filter design that accommodates both types of loops. The loop is implemented in digital hardware but with coefficients set in registers. In this way, the loop filter can be switched between perfect integrator and imperfect integrator for different phases of a space mission.

6 Acknowledgment

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

References

- [1] S. A. Stephens and J. B. Thomas, "Controlled-Root Formulation for Digital Phase-Locked Loops," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 31, No. 1, pp. 78-95, January 1995.
- [2] J. H. Yuen, ed., *Deep Space Telecommunications Systems Engineering*, Plenum Press, New York, 1983.
- [3] R. M. Jaffe and E. Rechtin, "Design and Performance of Phase-Lock Circuits Capable of Near Optimum Performance Over a Wide Range of Input Signals and Noise Levels," *IRE Transactions on Information Theory*, Vol. IT-1, pp. 67-76, March 1955.
- [4] A. J. Viterbi, *Principles of Coherent Communication*, McGraw-Hill, New York, 1966.